

Learning and Free Energies for Vector Approximate Message Passing

Alyson K. Fletcher and Philip Schniter

Abstract—Vector approximate message passing (VAMP) is a computationally simple approach to the recovery of a signal \mathbf{x} from noisy linear measurements $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{w}$. Like the AMP proposed by Donoho, Maleki, and Montanari in 2009, VAMP is characterized by a rigorous state evolution (SE) that holds under certain large random matrices and that matches the replica prediction of optimality. But while AMP's SE holds only for large i.i.d. sub-Gaussian \mathbf{A} , VAMP's SE holds under the much larger class: right-rotationally invariant \mathbf{A} . To run VAMP, however, one must specify the statistical parameters of the signal and noise. This work combines VAMP with Expectation-Maximization to yield an algorithm, EM-VAMP, that can jointly recover \mathbf{x} while learning those statistical parameters. The fixed points of the proposed EM-VAMP algorithm are shown to be stationary points of a certain constrained free-energy, providing a variational interpretation of the algorithm. Numerical simulations show that EM-VAMP is robust to highly ill-conditioned \mathbf{A} with performance nearly matching oracle-parameter VAMP.

I. INTRODUCTION

Consider the problem of estimating a random vector \mathbf{x} from linear measurements \mathbf{y} of the form

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{w}, \quad \mathbf{w} \sim \mathcal{N}(\mathbf{0}, \theta_2^{-1}\mathbf{I}), \quad \mathbf{x} \sim p(\mathbf{x}|\boldsymbol{\theta}_1), \quad (1)$$

where $\mathbf{A} \in \mathbb{R}^{M \times N}$ is a known matrix, $p(\mathbf{x}|\boldsymbol{\theta}_1)$ is a density on \mathbf{x} with parameters $\boldsymbol{\theta}_1$, \mathbf{w} is additive white Gaussian noise (AWGN) independent of \mathbf{x} , and $\theta_2 > 0$ is the noise precision (inverse variance). The goal is to estimate \mathbf{x} along while simultaneously learning the unknown parameters $\boldsymbol{\theta} := (\boldsymbol{\theta}_1, \theta_2)$ from the data \mathbf{y} and \mathbf{A} . This problem arises in Bayesian forms of linear inverse problems in signal processing, as well as in linear regression in statistics.

Even when the parameters $\boldsymbol{\theta}$ are known, exact estimation or inference of the vector \mathbf{x} is intractable for general priors $p(\mathbf{x}|\boldsymbol{\theta}_1)$. The *approximate message passing* (AMP) algorithm [1] and its generalization [2] are powerful, relatively recent, algorithms that iteratively attempt to recover \mathbf{x} . These methods are computationally fast and have been successfully applied to a wide range of problems, e.g., [3]–[11]. Most importantly, for large, i.i.d., sub-Gaussian random matrices \mathbf{A} , their performance can be exactly predicted by a scalar *state evolution* (SE) [12], [13] that provides testable conditions for optimality, even for non-convex priors. When the parameters $\boldsymbol{\theta}$ in the model are unknown, AMP can be combined with expectation

maximization (EM) methods [14]–[16] for joint estimation and learning.

As it turns out, the AMP methods [1], [2] are fragile with regard to the choice of the matrix \mathbf{A} , and can perform poorly outside the special case of zero-mean, i.i.d., sub-Gaussian \mathbf{A} . For example, AMP diverges with even mildly non-zero-mean and/or mildly ill-conditioned \mathbf{A} [17]. Several techniques have been proposed to improve the robustness of AMP including damping [17], [18], mean-removal [17], and sequential updating [19], but these remedies have limited effect.

Recently, the Vector AMP (VAMP) algorithm [20] was established as an alternative to AMP that is much more robust to the choice of matrix \mathbf{A} . In particular, VAMP has a rigorous SE that holds under large right-rotationally invariant \mathbf{A} , i.e., \mathbf{A} whose right singular-vector matrix is uniformly distributed on the group of orthogonal matrices. VAMP can be derived in several ways, such as through expectation propagation (EP) [21] approximations of belief propagation [20] or through expectation consistent (EC) approximation [22]–[24]. But the existence of a rigorous state evolution establishes it firmly in the class of AMP algorithms.

However, a shortcoming of the VAMP method [20] is that it requires that the parameters $\boldsymbol{\theta}$ in the model (1) are known. In this paper, we extend the VAMP method to enable learning of the parameters $\boldsymbol{\theta}$ via Expectation-Maximization (EM) [25], [26]. We call the proposed method EM-VAMP. As described below, exact implementation of EM requires estimating the posterior density $p(\mathbf{x}|\mathbf{y}, \boldsymbol{\theta})$ for each parameter estimate $\hat{\boldsymbol{\theta}}$. This is computationally not possible for the model (1). EM-VAMP is instead derived using a technique from Heskes [27] for combining EM with approximate inference of the posterior. Specifically, it is well-known that EM can be interpreted as a method to minimize a certain energy function. Here, we construct an approximation of the EM cost function that we call the EM-VAMP energy function and derive an algorithm to minimize this function.

Our main theoretical result shows that the fixed points of the EM-VAMP method are local minima of the EM-VAMP energy function and thus provide estimates of the parameters $\boldsymbol{\theta}$ and posterior density with a precise variational interpretation. By including the parameter learning, this result generalizes the fixed-point energy-function interpretation of EC given in [28], [29] and its variants [24].

Unfortunately, our results do not guarantee the convergence of the method to the fixed point. However, in numerical experiments on sparse regression problems, we show that the proposed method exhibits extremely stable convergence over a large class of matrices that cause AMP to diverge. Moreover,

A. K. Fletcher (email: akfletcher@ucla.edu) is with the Departments of Statistics, Mathematics, and Electrical Engineering, University of California, Los Angeles, CA, 90095. The work of A. K. Fletcher was supported by the NSF under grant CCF-1254204.

P. Schniter (email: schniter.1@osu.edu) is with the Department of Electrical and Computer Engineering, The Ohio State University, Columbus, OH, 43210. The work of P. Schniter was supported by the NSF under grant CCF-1527162.

the performance of EM-VAMP is almost identical to that of VAMP with known parameters. In particular, the method is able to obtain close to the theoretically optimal performance predicted by the replica method [30].

II. EM-VAMP

A. Review of VAMP

To describe the VAMP method in [20], we need to introduce some additional notation. First suppose that we can write the prior on \mathbf{x} as

$$p(\mathbf{x}|\boldsymbol{\theta}_1) = \frac{1}{Z_1(\boldsymbol{\theta}_1)} \exp[-f_1(\mathbf{x}|\boldsymbol{\theta}_1)], \quad (2)$$

where $f_1(\cdot)$ is some penalty function and $Z_1(\boldsymbol{\theta}_1)$ is a normalization constant. We assume that $f_1(\cdot)$ is separable, meaning that

$$f_1(\mathbf{x}|\boldsymbol{\theta}_1) = \sum_{n=1}^N f_{1n}(x_n|\boldsymbol{\theta}_1), \quad (3)$$

for scalar functions f_{1n} . This corresponds to the case that, conditional on $\boldsymbol{\theta}_1$, \mathbf{x} has independent components. Also, we write the likelihood for the Gaussian model (1) as

$$p(\mathbf{y}|\mathbf{x}, \boldsymbol{\theta}_2) := \frac{1}{Z_2(\boldsymbol{\theta}_2)} \exp[-f_2(\mathbf{x}, \mathbf{y}|\boldsymbol{\theta}_2)] \quad (4)$$

$$f_2(\mathbf{x}, \mathbf{y}|\boldsymbol{\theta}_2) := \frac{\boldsymbol{\theta}_2}{2} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|^2, \quad Z_2(\boldsymbol{\theta}_2) = \left(\frac{2\pi}{\boldsymbol{\theta}_2}\right)^{N/2}. \quad (5)$$

The joint density of \mathbf{x}, \mathbf{y} given parameters $\boldsymbol{\theta} = (\boldsymbol{\theta}_1, \boldsymbol{\theta}_2)$ is then

$$p(\mathbf{x}, \mathbf{y}|\boldsymbol{\theta}) = p(\mathbf{x}|\boldsymbol{\theta}_1)p(\mathbf{y}|\mathbf{x}, \boldsymbol{\theta}_2). \quad (6)$$

The VAMP algorithm [20] considers the case where the parameters $\boldsymbol{\theta}$ are known. In this case, VAMP attempts to compute *belief estimates* of the posterior density $p(\mathbf{x}|\mathbf{y}, \boldsymbol{\theta})$ of the form (for $i = 1, 2$)

$$b_i(\mathbf{x}|\mathbf{r}_i, \gamma_i, \boldsymbol{\theta}_i) \propto \exp\left[-f_i(\mathbf{x}, \mathbf{y}|\boldsymbol{\theta}_i) - \frac{\gamma_i}{2} \|\mathbf{x} - \mathbf{r}_i\|^2\right], \quad (7)$$

where the parameters \mathbf{r}_i, γ_i are optimized by the algorithm. To keep the notation symmetric, we have written $f_1(\mathbf{x}, \mathbf{y}|\boldsymbol{\theta}_1)$ for $f_1(\mathbf{x}|\boldsymbol{\theta}_1)$ even though the first penalty function does not depend on \mathbf{y} .

The steps of VAMP are identical to those shown for proposed EM-VAMP in Algorithm 1, except that VAMP skips the parameter updates in lines 4 and 11. Instead, VAMP fixes $\hat{\boldsymbol{\theta}}_{ik}$ for all iterations k . In Algorithm 1, we have focused on the MMSE version of VAMP since we are interested in approximate inference. There we use

$$\mathbb{E}[\phi(\mathbf{x})|\mathbf{r}_i, \gamma_i, \boldsymbol{\theta}_i] := \int \phi(\mathbf{x}) b_i(\mathbf{x}|\mathbf{r}_i, \gamma_i, \boldsymbol{\theta}_i) d\mathbf{x}$$

to denote the expectation with respect to the belief estimate $b_i(\cdot)$ in (7). Similarly, $\text{Cov}(\cdot|\cdot)$ is the covariance matrix with respect to the belief estimate and $\text{tr Cov}(\cdot|\cdot)$ is its trace. Hence, the VAMP method reduces the inference problem on the joint density (6) to computing expectations and variances with respect to the belief estimates (7).

One of the main motivations of the VAMP method is that, for the penalty functions (3) and (5) considered here, the

Algorithm 1 EM-VAMP

Require: Matrix $\mathbf{A} \in \mathbb{R}^{M \times N}$, penalty functions $f_i(\mathbf{x}, \mathbf{y}|\boldsymbol{\theta}_i)$, measurement vector \mathbf{y} , and number of iterations N_{it} .

```

1: Select initial  $\mathbf{r}_{10}, \gamma_{10} \geq 0, \hat{\boldsymbol{\theta}}_{1,-1}, \hat{\boldsymbol{\theta}}_{2,-1}$ .
2: for  $k = 0, 1, \dots, N_{\text{it}} - 1$  do
3:   // Input Denoising
4:    $\hat{\boldsymbol{\theta}}_{1k} = \arg \max_{\boldsymbol{\theta}_1} \mathbb{E}[\ln p(\mathbf{x}|\boldsymbol{\theta}_1)|\mathbf{r}_{1k}, \gamma_{1k}, \hat{\boldsymbol{\theta}}_{1,k-1}]$ 
5:    $\eta_{1k}^{-1} = (1/N) \text{tr}[\text{Cov}(\mathbf{x}|\mathbf{r}_{1k}, \gamma_{1k}, \hat{\boldsymbol{\theta}}_{1k})]$ 
6:    $\hat{\mathbf{x}}_{1k} = \mathbb{E}(\mathbf{x}|\mathbf{r}_{1k}, \gamma_{1k}, \hat{\boldsymbol{\theta}}_{1k})$ 
7:    $\gamma_{2k} = \eta_{1k} - \gamma_{1k}$ 
8:    $\mathbf{r}_{2k} = (\eta_{1k} \hat{\mathbf{x}}_{1k} - \gamma_{1k} \mathbf{r}_{1k}) / \gamma_{2k}$ 
9:
10:  // LMMSE estimation
11:   $\hat{\boldsymbol{\theta}}_{2k} = \arg \max_{\boldsymbol{\theta}_2} \mathbb{E}[\ln p(\mathbf{x}, \mathbf{y}|\boldsymbol{\theta}_2)|\mathbf{r}_{2k}, \gamma_{2k}, \hat{\boldsymbol{\theta}}_{2,k-1}]$ 
12:   $\eta_{2k}^{-1} = (1/N) \text{tr}[\text{Cov}(\mathbf{x}|\mathbf{r}_{2k}, \gamma_{2k}, \hat{\boldsymbol{\theta}}_{2k})]$ 
13:   $\hat{\mathbf{x}}_{2k} = \mathbb{E}(\mathbf{x}|\mathbf{r}_{2k}, \gamma_{2k}, \hat{\boldsymbol{\theta}}_{2k})$ 
14:   $\gamma_{1,k+1} = \eta_{2k} - \gamma_{2k}$ 
15:   $\mathbf{r}_{1,k+1} = (\eta_{2k} \hat{\mathbf{x}}_{2k} - \gamma_{2k} \mathbf{r}_{2k}) / \gamma_{1,k+1}$ 
16: end for
```

expectation and variance computations may be tractable at high dimensions. To understand why, first observe that, under the assumption of a separable penalty function (3), the belief estimate $b_1(\cdot)$ separates as

$$b_1(\mathbf{x}|\mathbf{r}_1, \gamma_1, \boldsymbol{\theta}_1) \propto \prod_{n=1}^N \exp\left[-f_1(x_n|\boldsymbol{\theta}_1) - \frac{\gamma_1}{2} (x_n - r_{1n})^2\right].$$

Thus, the expectation and variance computations in lines 5 and 6 decouple into N scalar computations. Furthermore, for the quadratic penalty (5), the belief estimate $b_2(\cdot)$ is Gaussian, i.e.,

$$b_2(\mathbf{x}|\mathbf{r}_2, \gamma_2, \boldsymbol{\theta}_2) \propto \exp\left[-\frac{\boldsymbol{\theta}_2}{2} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|^2 - \frac{\gamma_2}{2} \|\mathbf{r} - \mathbf{x}\|^2\right],$$

with mean and covariance given by

$$\mathbb{E}[\mathbf{x}|\mathbf{r}_2, \gamma_2, \boldsymbol{\theta}_2] = \mathbf{Q}^{-1} (\boldsymbol{\theta}_2 \mathbf{A}^T \mathbf{y} + \gamma_2 \mathbf{r}) \quad (8)$$

$$\text{Cov}[\mathbf{x}|\mathbf{r}_2, \gamma_2, \boldsymbol{\theta}_2] = \mathbf{Q}^{-1} \quad (9)$$

$$\mathbf{Q} = \boldsymbol{\theta}_2 \mathbf{A}^T \mathbf{A} + \gamma_2 \mathbf{I}. \quad (10)$$

Although (8)-(9) may suggest that VAMP requires an $N \times N$ matrix inverse at each iteration, it is shown in [20] that two $M \times N$ matrix-vector multiplications per iteration are sufficient if the SVD of \mathbf{A} is precomputed before initialization. Thus, VAMP reduces the intractable posterior inference problem to an iteration of N scalar estimation problems and 2 matrix-vector multiplies per iteration, just like AMP.

B. Learning the parameters $\boldsymbol{\theta}$

To learn the parameters $\boldsymbol{\theta}$, the EM-VAMP methods adds two steps, lines 4 and 11, to update $\hat{\boldsymbol{\theta}}_{ik}$. These maximizations are similar to those in the EM method, and we formalize this connection in the next section. The updates may be performed once per VAMP iteration, as written, or several times per VAMP iteration, since in practice this seems to speed convergence of EM-VAMP. For now, observe that due to the

structure of the prior in (2) and the likelihood in (4), we have that

$$\hat{\theta}_{i,k+1} = \arg \min_{\theta_i} \left\{ \mathbb{E} \left[f_i(\mathbf{x}, \mathbf{y} | \theta_i) \mid \mathbf{r}_{ik}, \gamma_{ik}, \hat{\theta}_{ik} \right] + \ln Z_i(\theta_i) \right\}. \quad (11)$$

This minimization is often tractable. For example, when the penalty function corresponds to an exponential family (i.e., $f_i(\mathbf{x}, \mathbf{y} | \theta_i) = \theta_i^\top \phi_i(\mathbf{x}, \mathbf{y})$ for sufficient statistic $\phi_i(\mathbf{x}, \mathbf{y})$), the minimization in (11) is convex. In particular, for the quadratic loss (5), the minimization is given by

$$\begin{aligned} \hat{\theta}_{2,k+1}^{-1} &= \frac{1}{N} \mathbb{E} [\|\mathbf{y} - \mathbf{A}\mathbf{x}\|^2 | \mathbf{r}_{2k}, \gamma_{2k}, \theta_{2k}] \\ &= \frac{1}{N} [\|\mathbf{y} - \mathbf{A}\mathbf{r}_{2k}\|^2 + \text{tr}(\mathbf{A}\mathbf{Q}_k^{-1}\mathbf{A}^\top)], \end{aligned} \quad (12)$$

where $\mathbf{Q}_k = \hat{\theta}_{2k}\mathbf{A}^\top\mathbf{A} + \gamma_{2k}\mathbf{I}$. As mentioned earlier, it is possible to reduce the complexity of evaluating (12) by precomputing the SVD of \mathbf{A} [20], since $\text{tr}(\mathbf{A}\mathbf{Q}_k^{-1}\mathbf{A}^\top) = \sum_{i=1}^R s_i^2 / (\theta_{2k}s_i^2 + \gamma_{2k})$ where $\{s_i\}_{i=1}^R$ are the non-zero singular values of \mathbf{A} . In this case, the update of θ_2 is very simple, computationally.

III. FIXED POINTS OF EM-VAMP

We will now show that the parameter updates in EM-VAMP can be understood as an approximation of the EM algorithm. We first briefly review the standard energy-function interpretation of EM [26]. Consider the problem of finding the maximum likelihood (ML) estimate of the parameter θ :

$$\hat{\theta} = \arg \max_{\theta} p(\mathbf{y} | \theta) = \arg \max_{\theta} \int p(\mathbf{x}, \mathbf{y} | \theta) d\mathbf{x}. \quad (13)$$

Due to the integration, this minimization is generally intractable. EM thus considers an auxiliary function,

$$Q(\theta, b) = -\ln p(\mathbf{y} | \theta) + D(b \| p(\cdot | \mathbf{y}, \theta)), \quad (14)$$

defined for an arbitrary density $b(\mathbf{x})$. In (14), $D(b \| p(\cdot | \mathbf{y}, \theta))$ is the KL divergence between $b(\mathbf{x})$ and the posterior density $p(\mathbf{x} | \mathbf{y}, \theta)$. Note that, for any parameter estimate θ ,

$$\min_b Q(\theta, b) = -\ln p(\mathbf{y} | \theta),$$

where the minimum occurs at the posterior $\hat{b}(\mathbf{x}) = p(\mathbf{x} | \mathbf{y}, \theta)$. Hence, the MLE (13) can, in principle, be found from the joint minimization

$$\hat{\theta} = \arg \min_{\theta} \min_b Q(\theta, b). \quad (15)$$

This fact leads to a natural alternating minimization,

$$\text{E-step: } \hat{b}_k = \arg \min_b Q(\hat{\theta}_k, b) = p(\mathbf{x} | \mathbf{y}, \hat{\theta}_k) \quad (16)$$

$$\text{M-step: } \hat{\theta}_{k+1} = \arg \min_{\theta} Q(\theta, \hat{b}_k). \quad (17)$$

This recursion is precisely the EM algorithm, written in a slightly non-standard form. Specifically, (16) is the E-step, which computes the posterior density of \mathbf{x} given \mathbf{y} and the current parameter estimate $\hat{\theta}_k$. A simple manipulation shows that

$$Q(\theta, b) = -\mathbb{E} [\ln p(\mathbf{x}, \mathbf{y} | \theta) | b] - H(b), \quad (18)$$

where the expectation is with respect to the density $b(\mathbf{x})$ and $H(b)$ is the differential entropy of b . Equation (18) shows that the minimization in (17) can equivalently be written as

$$\hat{\theta}_{k+1} = \arg \max_{\theta} \mathbb{E} [\ln p(\mathbf{x}, \mathbf{y} | \theta) | \hat{b}_k], \quad (19)$$

which is a familiar expression for the M-step. Unfortunately, the computation of the posterior density required by the E-step (16) is generally intractable for joint density (6) considered here.

We thus consider an alternate energy function, similar to that used by Heskes in [27] for understanding EM combined with belief propagation-based inference. First observe that, using (18) and (6), we can write the auxiliary function as

$$\begin{aligned} Q(\theta, b) &= \sum_{i=1}^2 \{ \mathbb{E} [f_i(\mathbf{x}, \mathbf{y} | \theta_i) | b] + \ln Z_i(\theta_i) \} - H(b) \\ &= \sum_{i=1}^2 D_i(b, \theta_i) + H(b), \end{aligned} \quad (20)$$

where $D_i(b, \theta_i)$ is the KL divergence,

$$D_i(b, \theta_i) = D \left(b \parallel Z_i(\theta_i)^{-1} e^{-f_i(\cdot | \mathbf{y} | \theta_i)} \right). \quad (21)$$

Now, given densities b_1, b_2 and q , we define the energy function

$$J(b_1, b_2, q, \theta) := D_1(b_1, \theta_1) + D_2(b_2, \theta_2) + H(q), \quad (22)$$

which matches the original auxiliary function $Q(\theta, b)$ under the matching condition $b = b_1 = b_2 = q$. Hence, we can rewrite the joint minimization (15) as

$$\hat{\theta} = \arg \min_{\theta} \min_{b_1, b_2} \max_q J(b_1, b_2, q, \theta) \text{ s.t. } b_1 = b_2 = q. \quad (23)$$

We call (22) the *EM-VAMP energy function*.

Now, as mentioned in the Introduction, VAMP—like many algorithms—can be viewed as an example of expectation consistent (EC) approximate inference [22]–[24]. Specifically, following the EC framework, we relax the above GFE optimization by replacing the constraints in (23) with so-called *moment matching* constraints:

$$\begin{aligned} \mathbb{E}(x_n | b_1) &= \mathbb{E}(x_n | b_2) = \mathbb{E}(x_n | q), \quad \forall n, \\ \mathbb{E}(\|\mathbf{x}\|^2 | b_1) &= \mathbb{E}(\|\mathbf{x}\|^2 | b_2) = \mathbb{E}(\|\mathbf{x}\|^2 | q). \end{aligned} \quad (24)$$

Thus, instead of requiring a perfect match in the densities b_1, b_2, q as in (23), we require only a match in their first moments and average second moments. Using the above approximation, we can then attempt to compute parameter estimates via the minimization

$$\hat{\theta} = \arg \min_{\theta} \min_{b_1, b_2} \max_q J(b_1, b_2, q, \theta) \text{ s.t. (24) are satisfied.} \quad (25)$$

Our main result shows that the fixed points of EM-VAMP are stationary points of the optimization (25). To state the result, we write the Lagrangian of the constrained optimization




Fig. 1. For sparse linear regression, recovery NMSE versus condition number of \mathbf{A} . Also shown is the replica prediction of the MMSE.

(25) as

$$L(b_1, b_2, q, \boldsymbol{\theta}, \boldsymbol{\beta}, \boldsymbol{\gamma}) := J(b_1, b_2, q, \boldsymbol{\theta}) - \sum_{i=1}^2 \beta_i^T [\mathbb{E}(\mathbf{x}|b_i) - \mathbb{E}(\mathbf{x}|q)] \\ + \sum_{i=1}^2 \frac{\gamma_i}{2} [\mathbb{E}(\|\mathbf{x}\|^2|b_i) - \mathbb{E}(\|\mathbf{x}\|^2|q)], \quad (26)$$

where $\boldsymbol{\beta} = (\beta_1, \beta_2)$ and $\boldsymbol{\gamma} = (\gamma_1, \gamma_2)$ represent sets of dual parameters for the first- and second-order constraints. We then have the following.

Theorem 1. *At any fixed point of the EM-VAMP algorithm with $\gamma_1 + \gamma_2 > 0$, we have*

$$\eta_1 = \eta_2 = \eta := \gamma_1 + \gamma_2, \quad (27a)$$

$$\hat{\mathbf{x}}_1 = \hat{\mathbf{x}}_2 = \hat{\mathbf{x}} := (\gamma_1 \mathbf{r}_1 + \gamma_2 \mathbf{r}_2) / (\gamma_1 + \gamma_2). \quad (27b)$$

Also, let $\beta_i := \gamma_i \mathbf{r}_i$, let \hat{b}_i be the density

$$\hat{b}_i(\mathbf{x}) := b_i(\mathbf{x}|\mathbf{r}_i, \gamma_i, \hat{\boldsymbol{\theta}}_i), \quad (28)$$

where $b_i(\cdot)$ is given in (7) and let $\hat{q}(\mathbf{x})$ be the Gaussian density

$$\hat{q}(\mathbf{x}) \propto \exp \left[-\frac{\eta}{2} \|\mathbf{x} - \hat{\mathbf{x}}\|^2 \right]. \quad (29)$$

Then, \hat{b}_i , $\hat{\boldsymbol{\theta}}$, and \hat{q} are critical points of the Lagrangian (26) that satisfy the moment matching constraints (24).

The proof is given in Appendix A and is an adaptation of a similar result in [24] with the addition of the parameters $\boldsymbol{\theta}$. The consequence of this result is that, if the algorithm converges, then its limit points are local minima of the EM-VAMP energy minimization.

IV. NUMERICAL EXPERIMENTS

While the above analysis characterizes the fixed points of EM-VAMP, it does not provide any guarantees on the convergence of the algorithm to the fixed points. To study




Fig. 2. For sparse linear regression, recovery NMSE versus iteration for condition number 32 in (a) and condition number 3162 in (b).

the convergence and evaluate the algorithm's performance, we conducted a numerical experiment.

We considered sparse linear regression, where the goal is to recover the signal \mathbf{x} from measurements \mathbf{y} from (1) without knowing the signal parameters $\boldsymbol{\theta}_1$ or the noise precision $\theta_2 > 0$. For our experiment, we drew \mathbf{x} from an i.i.d. Bernoulli-Gaussian (i.e., spike and slab) prior,

$$p(x_n|\boldsymbol{\theta}_1) = (1 - \beta_x)\delta(x_n) + \beta_x \mathcal{N}(x_n; \mu_x, \tau_x), \quad (30)$$

where parameters $\boldsymbol{\theta}_1 = \{\beta_x, \mu_x, \tau_x\}$ represent the sparsity rate $\beta_x \in (0, 1]$, the active mean $\mu_x \in \mathbb{R}$, and the active variance $\tau_x > 0$. Following [17], we constructed $\mathbf{A} \in \mathbb{R}^{M \times N}$ from the singular value decomposition (SVD) $\mathbf{A} = \mathbf{U}\mathbf{S}\mathbf{V}^T$, whose orthogonal matrices \mathbf{U} and \mathbf{V} were drawn uniformly with respect to the Haar measure and whose singular values s_i were constructed as a geometric series, i.e., $s_i/s_{i-1} = \alpha \forall i > 1$, with α and s_1 chosen to achieve a desired condition number $s_1/s_{\min(M,N)}$ as well as $\|\mathbf{A}\|_F^2 = N$. It is shown in [17], [18] that standard AMP (and even damped AMP) diverges when the matrix \mathbf{A} has a sufficiently high condition number. Thus, this matrix-generation model provides an excellent test for the stability of AMP methods. Recovery performance was assessed using normalized mean-squared error (NMSE) $\|\hat{\mathbf{x}} - \mathbf{x}\|^2 / \|\mathbf{x}\|^2$ averaged over 100 independent draws of \mathbf{A} , \mathbf{x} , and \mathbf{w} .

Figure 1 shows NMSE versus condition number for sparse linear regression under $M = 512$, $N = 1024$, $\beta_x = 0.1$, $\mu_x = 0$, and (τ_x, θ_2) giving a signal-to-noise ratio of 40 dB. EM-VAMP was initialized with $\beta_x = (M/2)/N$, $\tau_x = \|\mathbf{y}\|^2 / \|\mathbf{A}\|_F^2 \beta_x$, $\mu_x = 0$, and $\theta_2^{-1} = M^{-1} \|\mathbf{y}\|^2$. It is compared with (i) VAMP under perfect knowledge of $\boldsymbol{\theta} = \{\tau_w, \beta_x, \mu_x, \tau_x\}$; (ii) the EM-AMP algorithm from [15] with damping from [17]; and (iii) the replica prediction for Bayes minimum MSE from [31]. It was recently shown [32], [33] that the replica method gives the correct prediction in sparse linear regression when \mathbf{A} is i.i.d. Gaussian. Figure 1 shows that the NMSE of EM-VAMP is nearly indistinguishable from

that of VAMP and much more robust than EM-AMP to ill-conditioning in \mathbf{A} .

Figure 2(a) shows EM-VAMP and VAMP converging in ~ 10 iterations (whereas EM-AMP requires > 100 iterations) at condition number 32, and Figure 2(b) shows EM-VAMP converging in ~ 20 iterations at condition number 3162. These plots suggest that the convergence rate of EM-VAMP is i) nearly identical to that of genie-aided VAMP and ii) relatively insensitive to the condition number of \mathbf{A} . We note that, in generating the above figures, we used multiple updates of the noise precision θ_2 per VAMP iteration. In particular, (12) was iterated to convergence.

A Matlab implementation of our EM-VAMP method can be found in the GAMPmatlab software package at <http://sourceforge.net/projects/gampmatlab/>.

V. CONCLUSIONS AND FUTURE WORK

We presented an approach for recovering the signal \mathbf{x} from AWGN-corrupted linear measurements $\mathbf{y} = \mathbf{Ax} + \mathbf{w}$ by posing recovery in the MMSE framework while simultaneously learning the parameters $\boldsymbol{\theta}$ governing the signal prior $p(\mathbf{x}|\boldsymbol{\theta})$ and the AWGN variance. The proposed method combines EM and VAMP algorithms for approximate inference of the posterior. We showed that, if the algorithm converges, then its fixed points coincide with stationary points of a certain energy function. Simulations show the proposed method exhibits robustness to the condition number of \mathbf{A} and MMSE closely matching that of the replica prediction under known $\boldsymbol{\theta}$.

While the algorithm has great potential, one outstanding issue is that its convergence has not been established. One possible solution is to extend the convergence proofs in [24] or the state evolution analysis of VAMP [20]. Another avenue for future work is the application of EM-VAMP to *sparse Bayesian learning* (SBL) [34]. SBL tackles sparse linear regression using a Gaussian-scale-mixture prior $p(\mathbf{x}|\boldsymbol{\theta}_1) = \mathcal{N}(\mathbf{x}; \mathbf{0}, \text{Diag}(\boldsymbol{\theta}_1))$ with a deterministic known variance vector $\boldsymbol{\theta}_1 \in \mathbb{R}_+^N$ learned by the EM algorithm. While the standard SBL implementation uses an $N \times N$ matrix inverse at each EM iteration, the EM-VAMP implementation of SBL could avoid matrix inversions by precomputing an SVD.

REFERENCES

- [1] D. L. Donoho, A. Maleki, and A. Montanari, "Message-passing algorithms for compressed sensing," *Proc. Nat. Acad. Sci.*, vol. 106, no. 45, pp. 18914–18919, Nov. 2009.
- [2] S. Rangan, "Generalized approximate message passing for estimation with random linear mixing," in *Proc. IEEE Int. Symp. Inform. Theory*, Saint Petersburg, Russia, Jul.–Aug. 2011, pp. 2174–2178.
- [3] A. K. Fletcher, S. Rangan, L. Varshney, and A. Bhargava, "Neural reconstruction with approximate message passing (NeuRAMP)," in *Proc. Neural Information Process. Syst.*, Granada, Spain, Dec. 2011, pp. 2555–2563.
- [4] P. Schniter, "A message-passing receiver for BICM-OFDM over unknown clustered-sparse channels," *IEEE J. Sel. Topics Signal Process.*, vol. 5, no. 8, pp. 1462–1474, Dec. 2011.
- [5] S. Som and P. Schniter, "Compressive imaging using approximate message passing and a Markov-tree prior," *IEEE Trans. Signal Process.*, vol. 60, no. 7, pp. 3439–3448, Jul. 2012.
- [6] J. Ziniel and P. Schniter, "Dynamic compressive sensing of time-varying signals via approximate message passing," *IEEE Trans. Signal Process.*, vol. 61, no. 21, pp. 5270–5284, Nov. 2013.
- [7] J. P. Vila and P. Schniter, "An empirical-Bayes approach to recovering linearly constrained non-negative sparse signals," *IEEE Trans. Signal Process.*, vol. 62, no. 18, pp. 4689–4703, Sep. 2014.
- [8] U. S. Kamilov, V. K. Goyal, and S. Rangan, "Message-passing dequantization with applications to compressed sensing," *IEEE Trans. Signal Process.*, vol. 60, no. 12, pp. 6270–6281, Dec. 2012.
- [9] P. Schniter and S. Rangan, "Compressive phase retrieval via generalized approximate message passing," *IEEE Trans. Signal Process.*, vol. 63, no. 4, pp. 1043–1055, 2015.
- [10] J. Ziniel, P. Schniter, and P. Sederberg, "Binary linear classification and feature selection via generalized approximate message passing," *IEEE Trans. Signal Process.*, vol. 63, no. 8, pp. 2020–2032, 2015.
- [11] A. K. Fletcher and S. Rangan, "Scalable inference for neuronal connectivity from calcium imaging," in *Proc. Neural Information Processing Systems*, 2014, pp. 2843–2851.
- [12] M. Bayati and A. Montanari, "The dynamics of message passing on dense graphs, with applications to compressed sensing," *IEEE Trans. Inform. Theory*, vol. 57, no. 2, pp. 764–785, Feb. 2011.
- [13] A. Javanmard and A. Montanari, "State evolution for general approximate message passing algorithms, with applications to spatial coupling," *Information and Inference*, vol. 2, no. 2, pp. 115–144, 2013.
- [14] F. Krzakala, M. Mézard, F. Sausset, Y. Sun, and L. Zdeborová, "Statistical-physics-based reconstruction in compressed sensing," *Physical Review X*, vol. 2, no. 2, p. 021005, 2012.
- [15] J. P. Vila and P. Schniter, "Expectation-maximization Gaussian-mixture approximate message passing," *IEEE Trans. Signal Processing*, vol. 61, no. 19, pp. 4658–4672, 2013.
- [16] U. S. Kamilov, S. Rangan, A. K. Fletcher, and M. Unser, "Approximate message passing with consistent parameter estimation and applications to sparse learning," *IEEE Trans. Info. Theory*, vol. 60, no. 5, pp. 2969–2985, Apr. 2014.
- [17] J. Vila, P. Schniter, S. Rangan, F. Krzakala, and L. Zdeborová, "Adaptive damping and mean removal for the generalized approximate message passing algorithm," in *Proc. IEEE ICASSP*, 2015, pp. 2021–2025.
- [18] S. Rangan, P. Schniter, and A. Fletcher, "On the convergence of approximate message passing with arbitrary matrices," in *Proc. IEEE ISIT*, Jul. 2014, pp. 236–240.
- [19] A. Manoel, F. Krzakala, E. W. Tramel, and L. Zdeborová, "Swept approximate message passing for sparse estimation," in *Proc. ICML*, 2015, pp. 1123–1132.
- [20] S. Rangan, P. Schniter, and A. K. Fletcher, "Vector approximate message passing," *arXiv:1610.03082*, 2016.
- [21] T. P. Minka, "A family of algorithms for approximate Bayesian inference," Ph.D. dissertation, Dept. Comp. Sci. Eng., MIT, Cambridge, MA, 2001.
- [22] M. Opper and O. Winther, "Expectation consistent free energies for approximate inference," in *Proc. NIPS*, 2004, pp. 1001–1008.
- [23] —, "Expectation consistent approximate inference," *J. Mach. Learning Res.*, vol. 1, pp. 2177–2204, 2005.
- [24] A. K. Fletcher, M. Sahraee-Ardakan, S. Rangan, and P. Schniter, "Expectation consistent approximate inference: Generalizations and convergence," in *Proc. IEEE ISIT*, 2016, pp. 190–194.
- [25] A. Dempster, N. M. Laird, and D. B. Rubin, "Maximum-likelihood from incomplete data via the EM algorithm," *J. Roy. Statist. Soc.*, vol. 39, pp. 1–17, 1977.
- [26] R. Neal and G. Hinton, "A view of the EM algorithm that justifies incremental, sparse, and other variants," in *Learning in Graphical Models*, M. I. Jordan, Ed. MIT Press, 1998, pp. 355–368.
- [27] T. Heskes, O. Zoeter, and W. Wiegand, "Approximate expectation maximization," *NIPS*, vol. 16, pp. 353–360, 2004.
- [28] M. Opper and O. Winther, "Gaussian processes for classification: Mean-field algorithms," *Neural Computation*, vol. 12, no. 11, pp. 2655–2684, 2000.
- [29] —, "Adaptive and self-averaging Thouless-Anderson-Palmer mean-field theory for probabilistic modeling," *Physical Review E*, vol. 64, no. 5, p. 056131, 2001.
- [30] S. Rangan, A. Fletcher, and V. K. Goyal, "Asymptotic analysis of MAP estimation via the replica method and applications to compressed sensing," *IEEE Trans. Inform. Theory*, vol. 58, no. 3, pp. 1902–1923, Mar. 2012.
- [31] A. M. Tulino, G. Caire, S. Verdú, and S. Shamai, "Support recovery with sparsely sampled free random matrices," *IEEE Transactions on Information Theory*, vol. 59, no. 7, pp. 4243–4271, 2013.
- [32] J. Barbier, M. Dia, N. Macris, and F. Krzakala, "The mutual information in random linear estimation," *arXiv:1607.02335*, 2016.

- [33] G. Reeves and H. D. Pfister, "The replica-symmetric prediction for compressed sensing with Gaussian matrices is exact," in *Proc. IEEE ISIT*, 2016.
- [34] M. Tipping, "Sparse Bayesian learning and the relevance vector machine," *J. Machine Learning Research*, vol. 1, pp. 211–244, Sep. 2001.

APPENDIX

A. Proof of Theorem 1

The proof is modification of [24] with the addition of the parameters θ . From line 5 and 12 of Algorithm 1, $\eta_i = \gamma_1 + \gamma_2$ for $i = 1, 2$, which proves (27a). Also, since $\gamma_1 + \gamma_2 > 0$, we have that $\eta > 0$. In addition, from lines 8 and 15 we know,

$$\hat{\mathbf{x}}_i = (\gamma_1 \mathbf{r}_1 + \gamma_2 \mathbf{r}_2) / \eta \quad \text{for } i = 1, 2,$$

which proves (27b).

Now, by saying that $b_1, b_2, \hat{q}, \hat{\theta}$ are fixed points of the Lagrangian, we mean that

$$\hat{\theta} = \arg \min_{\theta} L(\hat{b}_1, \hat{b}_2, \hat{q}, \theta, \beta, \gamma), \quad (31)$$

$$(\hat{b}_1, \hat{b}_2) = \arg \min_{b_1, b_2} L(b_1, b_2, \hat{q}, \hat{\theta}, \beta, \gamma), \quad (32)$$

$$\hat{q} = \arg \max_q L(\hat{b}, q, \hat{\theta}, \beta, \gamma), \quad (33)$$

To prove (31), first observe that, for $i = 1, 2$,

$$\begin{aligned} L(b_1, b_2, q, \theta, \beta, \gamma) \\ &= J(b, q, \theta) + \text{const} \\ &= \mathbb{E}[f_i(\mathbf{x}, \mathbf{y}, \theta_i) | b_i] + \ln Z_i(\theta_i) + \text{const}, \end{aligned} \quad (34)$$

where the constant terms do not depend on θ and in the second step we used (21) and (22). Using (11) and the definition of $\hat{b}_i(\mathbf{x})$ in (32), we see that

$$\hat{\theta}_{i,k+1} = \arg \min_{\theta_i} \mathbb{E}[f_i(\mathbf{x}, \mathbf{y} | \theta_i) | \hat{b}_i] + \ln Z_i(\theta_i).$$

Combining this with (34) establishes (31).

To prove (32), we rewrite the Lagrangian (26) as

$$\begin{aligned} L(b, q, \theta, \beta, \gamma) \\ &\stackrel{(a)}{=} D(b_i \| e^{-f_i}) - \beta_i^T \mathbb{E}(\mathbf{x} | b_i) + \frac{\gamma_i}{2} \mathbb{E}[\|\mathbf{x}\|^2 | b_i] + \text{const} \\ &\stackrel{(b)}{=} D(b_i \| e^{-f_i}) + \frac{1}{2} \mathbb{E}[\gamma_i \|\mathbf{x} - \mathbf{r}_i\|^2 | b_i] + \text{const} \\ &\stackrel{(c)}{=} -H(b_i) + \mathbb{E}\left[f_i(\mathbf{x}, \mathbf{y}, \hat{\theta}_i) + \frac{\gamma_i}{2} \|\mathbf{x} - \mathbf{r}_i\|^2 \middle| b_i\right] + \text{const} \\ &\stackrel{(d)}{=} D(b_i \| \hat{b}_i) + \text{const}, \end{aligned} \quad (35)$$

where in step (a) we removed the terms that do not depend on b_i ; in step (b) we used the fact that $\beta_i = \gamma_i \mathbf{r}_i$; and in steps (c) and (d) we used the definitions of KL divergence and \hat{b}_i in (28). Thus, the minimization in (32) yields (28).

The maximization over q in (33) is computed similarly. Removing the terms that do not depend on q ,

$$\begin{aligned} L(b_1, b_2, q, \beta, \gamma) \\ &= H(q) + \sum_{i=1}^2 \beta_i^T \mathbb{E}(\mathbf{x} | q) - \frac{\gamma_i}{2} \sum_{i=1}^2 \mathbb{E}[\|\mathbf{x}\|^2 | b_i] + \text{const} \\ &\stackrel{(a)}{=} H(q) + \eta \hat{\mathbf{x}}^T \mathbb{E}(\mathbf{x} | q) - \frac{\eta}{2} \mathbb{E}[\|\mathbf{x}\|^2] + \text{const} \\ &\stackrel{(b)}{=} H(q) - \frac{\eta}{2} \mathbb{E}[\|\mathbf{x} - \hat{\mathbf{x}}\|^2] + \text{const} \\ &\stackrel{(c)}{=} -D(q \| \hat{q}) + \text{const}, \end{aligned} \quad (36)$$

where step (a) uses the facts that $\gamma_1 + \gamma_2 = \eta$ and

$$\beta_1 + \beta_2 = \gamma_1 \mathbf{r}_1 + \gamma_2 \mathbf{r}_2 = \eta \hat{\mathbf{x}},$$

step (b) follows by completing the square, and step (c) uses the density in (29). Hence, the maximizer of (33) is given by (29).

Also, from the updates of $\hat{\mathbf{x}}_i$ and η_i in Algorithm 1, we have

$$\hat{\mathbf{x}} = \mathbb{E}(\mathbf{x} | b_i), \quad \eta^{-1} = \frac{1}{N} \text{tr Cov}(\mathbf{x} | b_i).$$

Since \hat{q} is Gaussian, its mean and average covariance are

$$\mathbb{E}(\mathbf{x} | q) = \hat{\mathbf{x}}, \quad \frac{1}{N} \text{tr Cov}(\mathbf{x} | q) = \eta^{-1}.$$

This proves that the densities satisfy the moment matching constraints (24).